High Performance SQL Server with Storage Center 6.4 All Flash Array

Dell Storage
November 2013

A Dell Compellent Technical White Paper

Revisions

Date	Description
November 2013	Initial release

Table of contents

Acknowledgements

Dell Compellent Technical Solutions: Doug Bernhardt, SQL Server Product Specialist

Dell Storage Technical Marketing: Omar Rawashdeh and Magi Kapoor

Editing: Camille Daily

Executive summary

Online transaction processing (OLTP) systems, often referred to as transactional systems, are designed to process small, quick, interactive workloads where users expect a fast response time. These applications are optimal for managing rapidly changing data and typically have many users performing transactions while at the same time changing real-time data. Although individual data requests by users usually reference few records, many of these requests are being made at the same time.

These systems are at the core of business operations and their design and development become more challenging with large volumes of data, users, and transactions rates. Moreover, the demand for greater I/O performance is driven by new applications such as virtualization and cloud computing. As increased data load requirements continue to outpace traditional storage I/O technology, businesses are challenged with how to accelerate performance to meet the new increased transactional processing workloads. The emergence of flash storage technology offers much faster I/O performance.

The Dell Compellent flash optimized solution is a high performance and low latency solution for OLTP systems and other I/O intensive workloads. This Dell solution optimizes cost versus performance versus capacity tradeoffs by integrating write intensive (WI) Solid State Drives (SSD) with lower cost, higher capacity read intensive (RI) SSDs and rotating disks in a single solution.

This technical report describes a tested and validated solution for designing and deploying online transaction processing applications with Microsoft SQL Server 2012 using the Compellent SC8000 storage arrays with all flash technology.

The included study demonstrates the performance capability of the all flash array with Storage Center 6.4 and how it can be a key component in designing high performance SQL Server architectures for OLTP workloads. This paper also compares the all flash solution against spinning disk in terms of performance, footprint, and value.

1 Changing the economics of storage-flash at the price of spinning Disk

Dell provides an all flash or hybrid flash solution that is price competitive with other flash solutions on the market. The Dell Compellent innovative flash optimized solutions offer about twice the price advantage over competitive hybrid solutions and approximately five times the price advantage over competitive all flash solutions. The flash optimized solution is made possible by the introduction of RI SSDs and Storage Center 6.4 enhancements to data progression to support tiering across (WI) SSDs (Tier 1) and RI SSDs (Tier 2). Compellent flash optimized solutions are designed for I/O intensive applications (such as OLTP) that require random I/O per second (IOPS) and ultra-fast responsiveness.

This flash enclosure, designed to be used for all flash and hybrid flash arrays, supports multi-tiering of flash between higher capacity and lower cost RI SSDs, WI SSDs, and the traditional rotating drives (see Figure 1). Combined with flash enclosures, Storage Center 6.4 delivers high-performance flash storage at a fraction of the cost compared to other storage solutions.

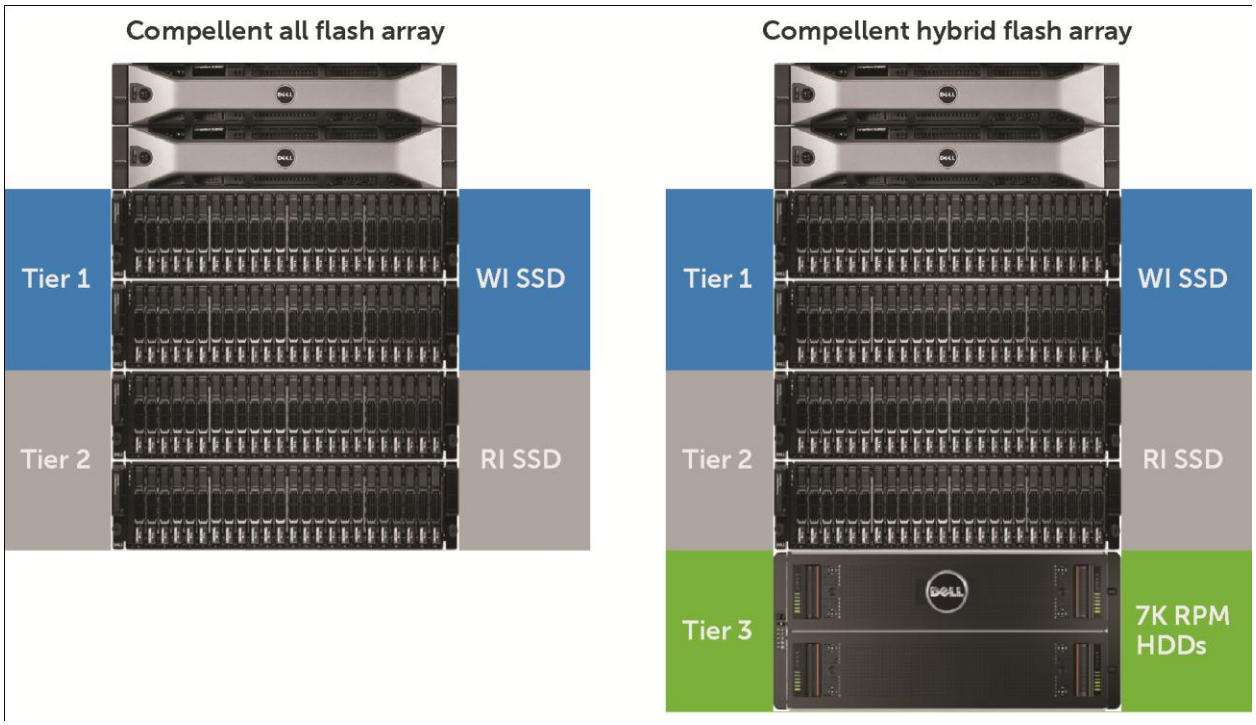

Figure 1 Compellent All Flash and Hybrid arrays

1.1 Cost savings

Organizations are increasingly lowering IT costs by optimizing processes and acquiring storage platforms that can do more with less. There is an acute need for low cost high performance SAN architectures that can keep up with the ever increasing workload demands. Enterprises using Microsoft SQL Server 2012 must consider new approaches to their performance challenges.

The Compellent flash-optimized storage solution needs less hardware than a traditional storage platform for the same workload (Figure 1). This lowers overall storage total cost of ownership (TCO) by minimizing:

- Power consumption and cooling: Up to 83% based on the tested configurations.
- Data center footprint: Up to 75% based on the tested configurations.
- Overall cost: Minimized cost for storage software licensing and support.

2 Test configuration and components

OLTP database applications are optimal for managing rapidly changing data. These applications typically have many users who are performing transactions and changing data simultaneously. Optimizing an OLTP database system running on SQL Server allows the maximum number of transactions through the system in the least amount of time.

The key metric in measuring performance of OLTP workloads is the number of IOPS that can be achieved while maintaining a healthy response time. For OLTP transactions to take place, SQL Server relies on an efficient I/O subsystem. According to a Microsoft SQL Server best practices article (http://technet.microsoft.com/en-us/library/cc966412.aspx), a SQL Server OLTP transaction profile is composed of the following pattern:

- OLTP processing is generally random in nature for both reads and writes issued against data files.
- Read activity (in most cases) is constant in nature.
- Write activity to the data files occur during checkpoint operations.
- Log writes are sequential in nature with a size that is dependent on the nature of the workload.
- Log reads are sequential in nature.

A test scenario was designed and configured to evaluate the performance of Storage Center 6.4 with all flash array compared to 15K spinning drives under a typical database workload. An OLTP type database workload was chosen to generate the database activity. The OLTP type database workload was chosen because it produces database activity that simulates real world usage patterns for the purpose of comparing system components.

2.1 System configuration

The same test scenario was executed on a single system hardware configuration. For each test a different storage configuration was used. The only differences in the storage configurations were the type of disks and the number of disks allocated to the Storage Center volumes. These changes were performed within the Storage Center user interface without changes to the hardware between tests.

The common components in the tests were the application server, the database server, and the Storage Center including all network components used for connectivity.

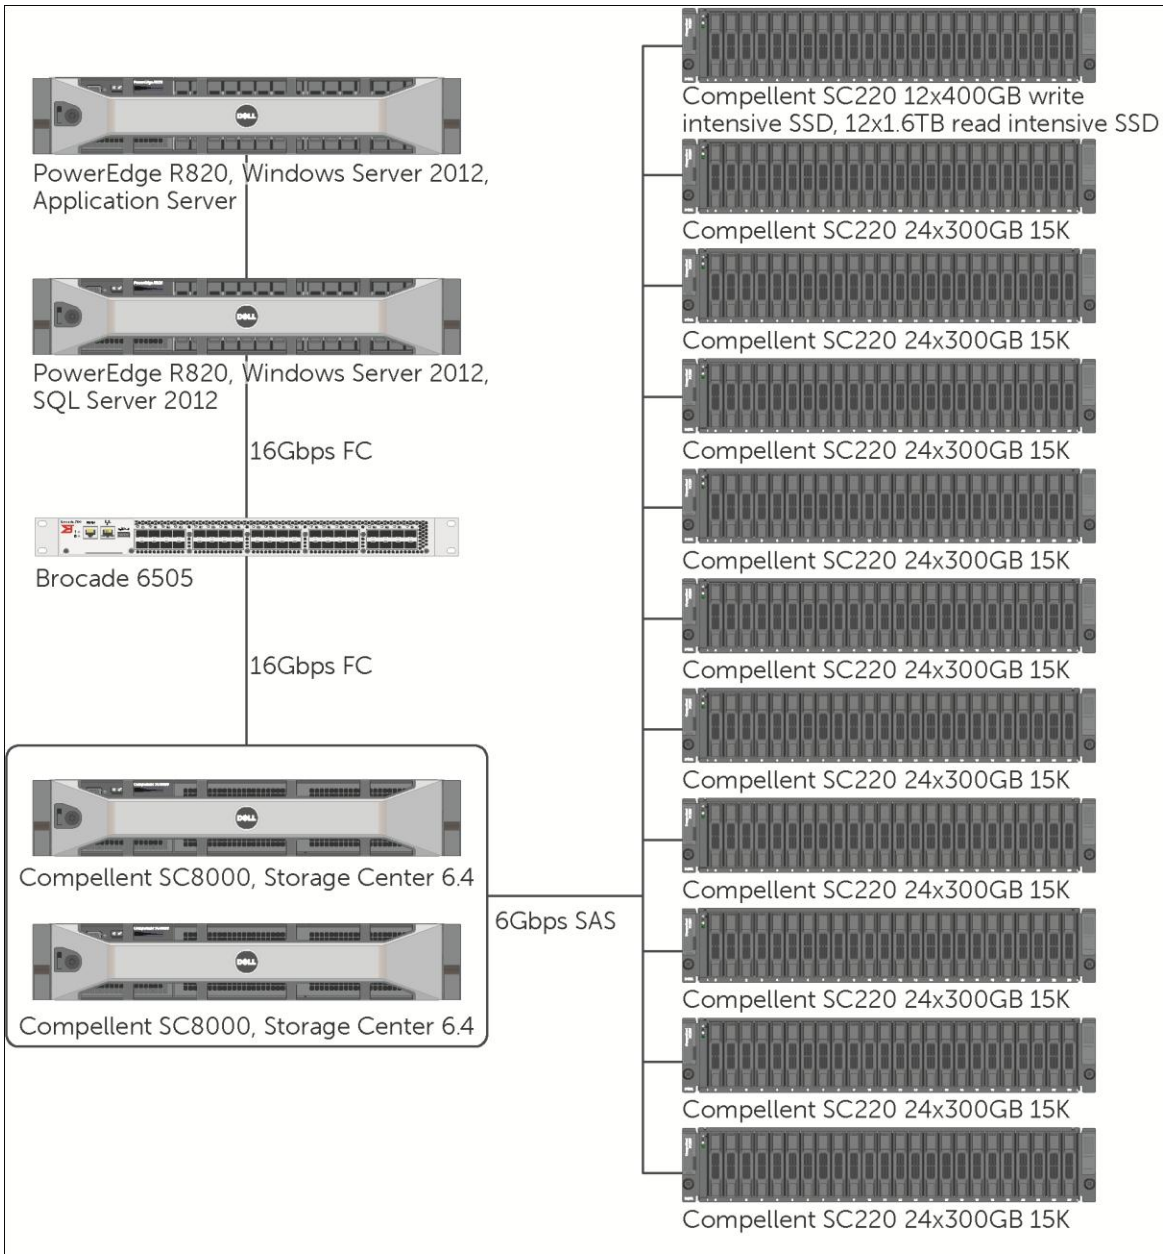

Figure 2 Test system configuration

Key design details of the test system configuration shown in Figure 2 include:

- A single R820 Dell PowerEdge server was used to host a single SQL Server instance running SQL Server 2012 on Windows Server 2012.
- The application server was running an OLTP type workload to produce the database workload.
- End to End 16Gbps Fibre Channel was used between the database server and the Storage Center. The Host Bus Adapters used in the Storage Center and the Database Server were QLogic QLE2662 adapters offered by Dell.
- The Dell Compellent Storage Center was configured with Dual SC8000 controllers running Storage Center 6.4 to drive three different disk configurations.
- A total of three SAS chains and 11 SC220 enclosures were used for the test. The all flash enclosure was configured on a separate SAS chain. Two additional SAS chains were configured with five enclosures of 24 300 GB 15K drives each.

2.2 Storage configuration

In order to understand the benefits of flash storage on OLTP workloads, three different storage configurations were tested. An OLTP type workload was run on the application server using a single SQL Server 2012 database. Two Dell Compellent SC8000 controllers were used in a dual controller configuration for the three storage configurations tested.

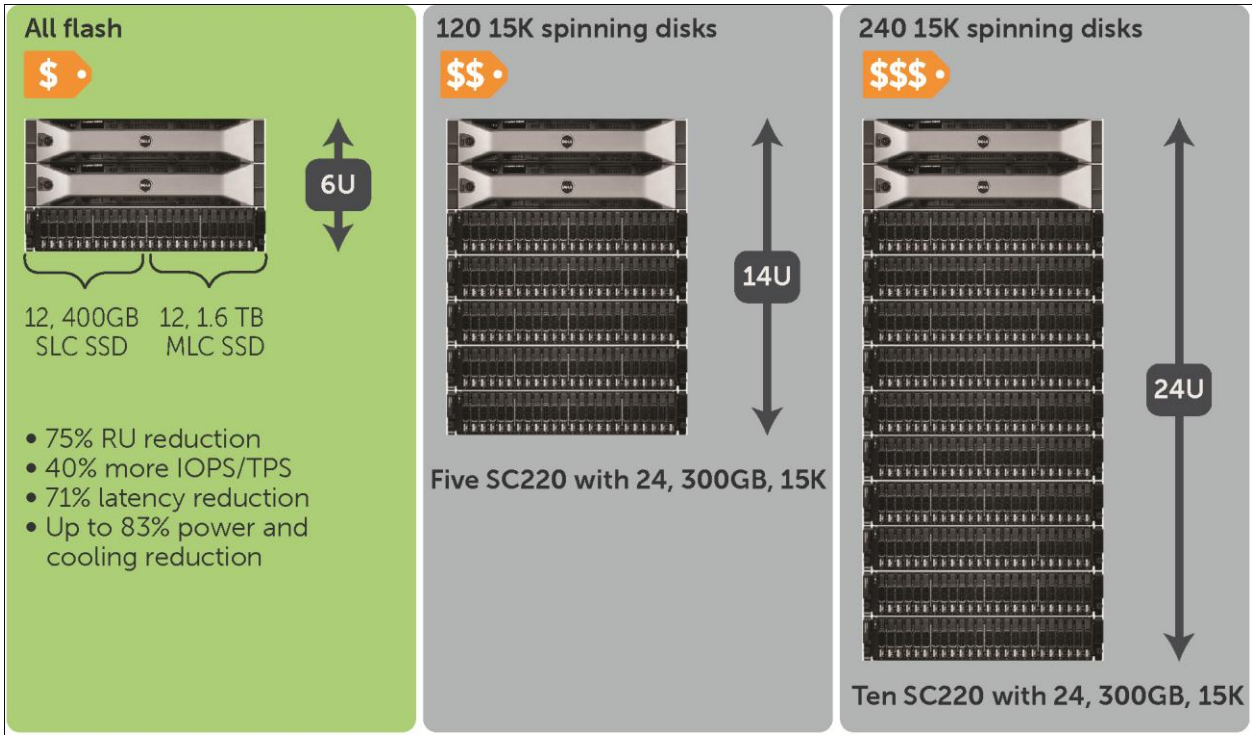

Figure 3 Tested system configurations

The first storage configuration (from left to right) that was tested was an all flash configuration of a single SC-220 enclosure containing twelve 400GB WI SSD drives and twelve 1.6TB RI SSD drives. This drive

configuration was chosen to illustrate a combination of performance, capacity, and storage density. The number of drives and mix of write and read intensive drives provides a good mix of performance and capacity. It is common to use full enclosures to maximize rack space and storage capacity.

The second storage configuration that was tested consisted of five SC220 enclosures with 24 300GB 15K spinning drives for a total of 120 15K drives. This spinning disk configuration was chosen because it is roughly the same cost as the all-flash configuration. The purpose is to demonstrate the performance that can be achieved for the same price on spinning disk vs. all-flash.

The final storage configuration that was tested consisted of ten SC220 enclosures with 24 300GB 15K spinning drives for a total of 240 15K drives. This spinning disk configuration was chosen to demonstrate the performance scalability compared to the all flash configuration.

2.3 Key test components

The key components and the reasons they were chosen for this performance study are listed in the following sections.

2.3.1 SQL Server 2012

As of November 2013, SQL Server 2012 is the current released version of SQL Server DBMS designed for high performance database applications. For these tests, SQL Server was configured to use 120GB of memory and the maximum degree of parallelism (MAXDOP) was set to 1. All other settings were left at the defaults. This configuration was used to emulate a real-world environment. There was no special tuning performed specifically for this type of test.

2.3.2 Dell PowerEdge R820 server

The Dell PowerEdge R820 server was chosen for this test since it provides outstanding processing power in a 2U space. The R820 is currently the most powerful server offered by Dell in a 2U space. There are several articles published on Dell.com that highlight the performance of the PowerEdge R820. (See the "Resources" section under the R820 product page for these articles)

2.3.3 Compellent SC8000 Storage Center

Compellent Storage Center version 6.4 was chosen due to its support for the all flash array configuration and the ability to allow data to automatically move from high performance WI/SLC SSD to high capacity RI/MLC SSD via Data Progression. Storage Center 6.4 also contains performance improvements for flash optimized environments as well as a gas gauge feature that indicates SSD life expectancy.

2.4 Windows Performance Monitor

Performance data was collected on the database server via Windows Performance Monitor (Perfmon). Perfmon is a very powerful performance monitoring tool that is included with Windows Server and can be customized to capture performance statistics. The performance data was sampled every five seconds and then averaged over a one hour test run preceded by a 20 minute ramp up period.

3 Test results

The results from tests run on three different storage configurations are reported in terms of application performance and storage performance. The application performance results indicate how customer applications would improve with an all flash configuration as compared to spinning disk. The storage performance results explain what was occurring with the tested configurations.

3.1 Application performance

Application performance is the experience of the customer or the business. Performance sensitive applications may be implemented to serve users as fast as possible or to perform work as quickly as possible. Application performance can also be used to compare systems to assess system consolidation, scalability, or other purposes.

The primary metric captured in this workload was transactions per second (TPS). This metric is important because it is a direct measure of the amount of work being performed in a given amount of time. The transactions used in this test are modeled after real-world transactions and provide a way to compare different systems. Comparing the TPS numbers provides a way to evaluate the performance of different systems. For the purposes of this paper, TPS measurements were used to compare the capabilities of tested storage configurations and demonstrate the performance of the Compellent all flash array,

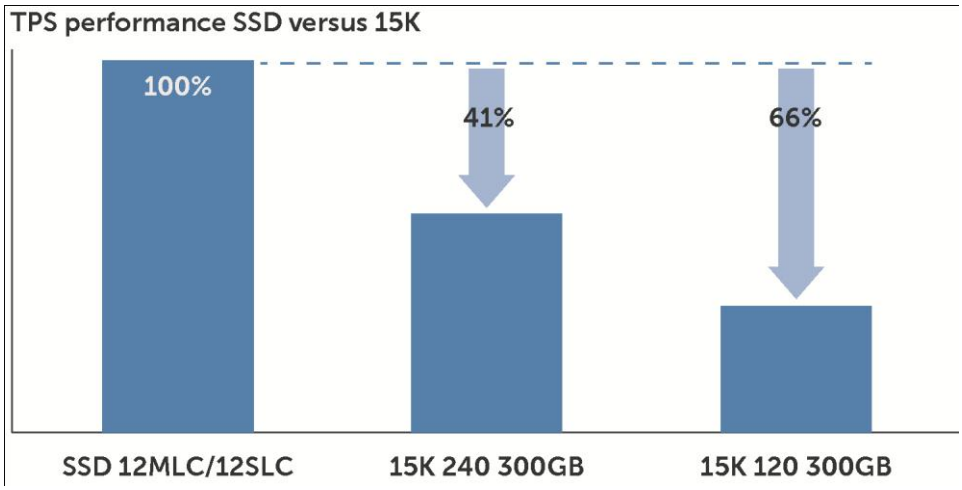

Figure 4 TPS performance: all flash versus spinning disk

When the all flash configuration is compared to the cost equivalent configuration of 120-15K Spinning disk, the all flash configuration produces three times the throughput. Doubling the number of 15K drives from 120 to 240 only yielded a 73% improvement in throughput. Some key points from the flash versus spinning disk comparison:

- OLTP application TPS increase significantly because storage latency is reduced. SQL Server is spending less time waiting on storage and therefore is able to process the next request faster.
- Since I/O latency is typically a large component of database transactions, reduced I/O latency will likely result in improved user response times.

3.2 Storage performance

High performing storage is critical in database applications, especially those that process OLTP transactions. Since the majority of time typically spent in OLTP transactions is accessing storage, the storage performance has a huge impact on database performance. The results in IOPS and latency categories explain the dramatic improvement in TPS performance. There was a large increase in the number of IOPS and they were completed much faster.

3.2.1 IOPS

The all flash configuration outperformed the cost equivalent spinning disk configuration by nearly three to one. Doubling the number of spinning drives from 120 to 240 only yielded a 68% improvement in the number of IOPS.

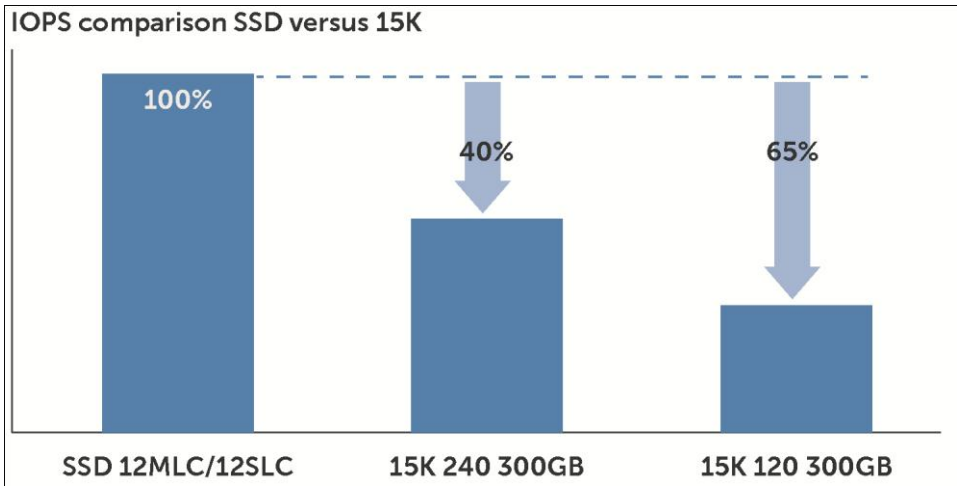

Figure 5 IOPS comparison

3.2.2 Latency

The low latency that SSD drives can provide is the key to the results in this test scenario. For the all flash configuration above, the average latency measured at the database server was 10 times lower than that of the 120 15K drives. Doubling the number of spinning drives resulted in improved latency by 64%. While this is an improvement, it is still two and half times higher than the all flash configuration. Due to the mechanical limits of the spinning drives, no matter how many are added, the latency will not match the SSD configuration.

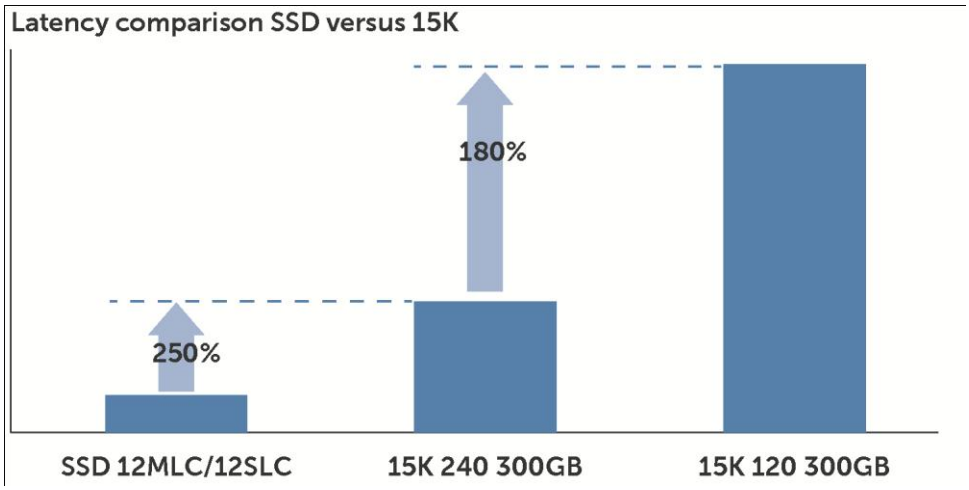

Figure 6 Latency comparison

4 Conclusion

The Dell Compellent all flash arrays have changed the way storage is designed. For the same cost, an all flash versus a spinning disk solution provides better performance. The keys to this performance are the low latency that flash drives provide and the flash-optimized performance of Storage Center 6.4. In addition to the outstanding performance, Storage Center 6.4 offers other benefits as well. The flexibility of combining read intensive and write intensive SSD offers the ability to offer both capacity and performance on the flash platform. The Compellent flash-optimized solution offers flash at the price of disk and provides an outstanding storage solution for running high performance SQL Server workloads.

A Test configuration specifications

A.1 Hardware specifications

Application Server	
Model	PowerEdge R820
CPU	4 - E5-4617
Memory	16 – 256GB DDR-3 1333MHz (256 GB Total)
Network	Integrated Broadcom Gigabit Ethernet BCM5720
Database Server	
Model	PowerEdge R820
CPU	4 - E5-4650
Memory	16 – 256GB DDR-3 1333MHz (256 GB Total)
Network	Integrated Broadcom Gigabit Ethernet BCM5720
Host Bus Adapter	2 - QLogic 2662 16Gbps
Fiber Channel Switch	
Model	Brocade 6505
Ports	24
Speed	16Gbps
Storage Controller	
Model	2 - SC8000
Host Bus Adapter	2 - QLogic 2662 16Gbps
IO Cards	LSI SAS 9201

A.2 Software and firmware versions

Vendor	Model	Version
Dell	Storage Center	6.4.1
Microsoft	Windows Server 2012 Standard Edition	6.2.9200
Microsoft	SQL Server 2012 Enterprise Edition	11.0.3000.0
Dell	QLE2662 Host Bus Adapter	6.06.03

A.3 Storage Center configuration

Test 1 – All Flash	
Disk Type	12-400 GB Write Intensive SSD, 12-1.6 TB Read Intensive SSD
Number of Enclosures	1
Hot Spares	2 – One for each type
Managed Disks	22
Test 2 – 120 15K Drives	
Disk Type	120-300 GB 15K
Number of Enclosures	5
Hot Spares	5 – One per enclosure
Managed Disks	115
Test 3 – 240 15K Drives	
Disk Type	240-300 GB 15K
Number of Enclosures	10
Hot Spares	10 – One per enclosure
Managed Disks	230

A.4 SQL Server configuration

SQL Server	
SQL DB Files- Volume size	R: Volume (Controller 1) Primary – 8MB 2 files – Fixed Filegroup – 10 MB each 2 files – Growing Filegroup – 768 GB each 2 files – Scaling Filegroup – 20 GB Each S: Volume (Controller 2) 2 files – Fixed Filegroup – 10 MB each 2 files – Growing Filegroup – 768 GB each 2 files – Scaling Filegroup – 20 GB Each T: Volume (Controller 1) 1 file –transaction log – 200 GB
Max Server Memory	120 GB
Degree of Parallelism	1
Database Workload	OLTP type workload
Database size	3 TB – 2TB data, 1TB free
Number of users (simulated)	575

B Additional resources

Support.dell.com is focused on meeting your needs with proven services and support.

DellTechCenter.com is an IT Community where you can connect with Dell Customers and Dell employees for the purpose of sharing knowledge, best practices, and information about Dell products and installations.

Use the whitepapers, tips, and best practice guides in the Dell Compellent Knowledge Center to help get the most out of your Dell Compellent Storage Center. The Knowledge Center can be accessed at http://kc.compellent.com.

Referenced or recommended Microsoft publications:

- Microsoft SQL Server 2008: Disk Partition Alignment Best Practices for SQL Server
  http://msdn.microsoft.com/en-us/library/dd758814.aspx

- SQL Server I/O Basics
  http://technet.microsoft.com/en-us/library/cc966500.aspx

- Pre-deployment I/O Best Practices
  http://technet.microsoft.com/en-us/library/cc966412.aspx

- Analyzing I/O Characteristics and Sizing Storage Systems for SQL Server Database Applications
  http://msdn.microsoft.com/en-us/library/ee410782(v=SQL.100).aspx

- Microsoft SQL Server homepage
  http://www.microsoft.com/sql

- SQL Server Customer Advisory Team
  http://sqlcat.com